

(12) UK Patent Application (19) GB (11) 2 358 100 (13) A

(43) Date of A Publication 11.07.2001

(21) Application No 0029063.5

(22) Date of Filing 29.11.2000

(30) Priority Data

(31) 09449618

(32) 30.11.1999

(33) US

(71) Applicant(s)

International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America

(72) Inventor(s)

Gordon Wesley Braudaway
Marco Martens
James B Shearer
Charles Philippe Louis Tresser
Cahi W Wu

(74) Agent and/or Address for Service

S R Davies
IBM United Kingdom Limited, Intellectual Property
Dept, Hursley Park, WINCHESTER, Hampshire,
SO21 2JN, United Kingdom

(51) INT CL⁷

H04N 1/32 // G06T 1/00

(52) UK CL (Edition S)

H4F FBB F13A F22

(56) Documents Cited

EP 0933919 A2 WO 99/10858 A2 WO 97/34391 A1
IEEE trans. Multimedia, Vol 2 No 4, Dec 2000, C-S Lu
ET AL, "Cocktail Watermarking...", pp 209-224

(58) Field of Search

UK CL (Edition S) H4F FBB
INT CL⁷ G06T 1/00 , H04N 1/00 1/32 5/913
ONLINE: WPI; JAPIO; EPODOC; INSPEC; IEL

(54) Abstract Title

Digital watermarks

(57) The present invention relates to a digital watermarking method and system which encode different pairs of watermarks into each of a plurality of images offered for use by a vendor. The watermarks in each pair are derived from two separate collections (COL1 and COL2) of watermarks and sufficiently different so as to prevent false positives. Because each pair of watermarks is assigned to a different customer relative to a particular image, unauthorised use of a digital image sold to a customer may be determined by locating the associated pair of watermarks assigned to the customer in the image. Collusion detection is also realized by forming each pair of masks from sub-collections of masks which are detectable in an image formed by combining the same images sold to one or more customers.

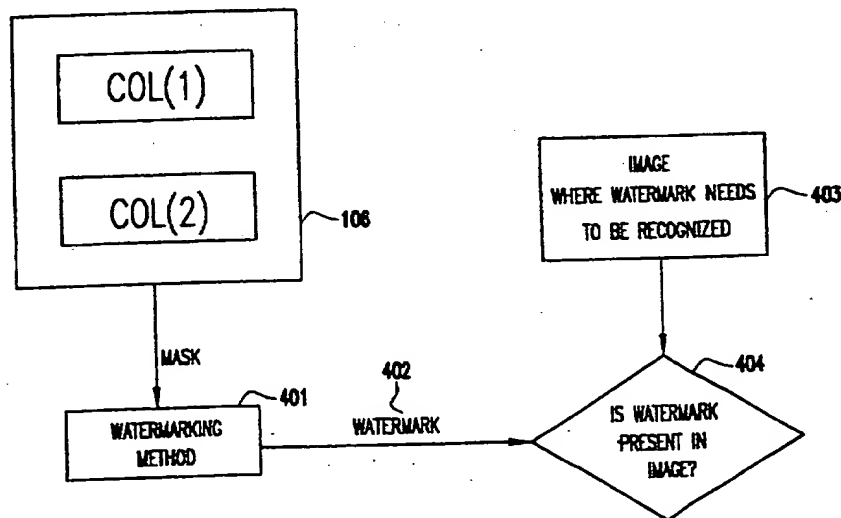


FIG.4

GB 2 358 100 A




FIG.1




FIG. 2




FIG.3




FIG.4

DIGITAL WATERMARKS

The present invention generally relates to digital watermarks, and, more particularly, to a digital watermarking method and system that facilitate digital image identification.

An imperceptible watermark is defined as an alteration of a data set which is, for the most part, imperceptible to a human, i.e., the watermark should be invisible or almost invisible, but can be recognised by a machine such as a computer. The general principle of developing a watermark is disclosed, for example, by M. M. Yeung et al. in "Digital Watermarking for High-Quality Imaging", *Proceedings of the IEEE Signal Processing Multimedia Workshop*, Princeton, New Jersey (1997).

A robust watermark is one conceived to survive modifications of the image, and more precisely, to still be detectable when the image has been modified to some reasonable extent. Robust watermarks are usually required to be non-removable by an adverse party, at least not removable without visible distortion to the image. They are usually designed, for instance, to establish ownership or to help protect copyrights. In several cases, one expects them to be detectable when the image is transferred to the analog world, such as when printed or transformed into analog signals such as required for display on a TV screen.

Watermarks cannot be replaced by digital signatures in this context, as known signature schemes with no visible trace on the image do not resist transfer from the digital to analog world. Due to the dual constraint of invisibility and resistance to minor changes in the image, robust watermark detection is therefore based on statistical analysis: any individual components of a mark can be altered by attack and/or by digital/analog (or D/A) conversion so that one is reduced to check that traces of the overall mark persist.

In general, a robust watermark consists of an array $M(K)$ of elements, $M(K)(h,v)$ defined by a cryptographically designed key K . Here, h and v represent horizontal and vertical coordinates. According to one technique for generating a robust watermark, $M(K)$ is a matrix of "0's" and "1's", where a zero means a pixel gets brighter, and a 1 in the matrix means a pixel gets darker after watermarking (one can have a richer set of possibilities where, for instance, a "2" in the matrix means that no change is made.) The watermark $M(K)$ (referred to as a "mask") is used to modify some pixels attributes according to some algorithm, defining some function

F , which generates the modified pixel attribute $m(h,v)$ at location (h,v) as a function of the mask M , the pixel's attribute before marking $A(h,v)$, and the attributes before marking $A_n(h,v)$ of neighboring pixels, i.e.,

$$m(h,v) = F(M(K)(h,v), A(h,v), A_n(h,v)).$$

5 The amount by which watermarking modifies the luminosity at a given pixel is not necessarily uniform for better visual performance and/or for better resistance to attacks. As the image and the accompanying watermark are expected to suffer some modifications, the key K , as well as the function F , are usually chosen as secret keys: a description of
10 cryptographic techniques, with directions on how to use several of their implementations, can be found in *Handbook of Applied Cryptography* by Alfred J. Menezes, Paul C. van Oorschot and Scott A. Vanstone, CRC Press, 1997.

 An example of a robust invisible robust watermark is disclosed, for example, in US Patent No. 5,825,892 issued to Mintzer and Braudaway.

15 Digital watermarking is directed, at least in part, to addressing the following business case. Suppose an image vendor (called The Ven Company in this hypothetical situation) sells digital images to a pool of customers C_1, C_2, \dots, C_N . Not all images are sold necessarily to all customers. Each customer indeed only buys the right to print these images in some
20 publications. No customer has the right to circulate or sell the digital form of the images.

 The Ven Company wants to prevent unauthorised resale of these images. In particular, if one of the images is found in some publication not listed by any customer, The Ven Company will legitimately suspect that one of its
25 customers has broken his contract. It would therefore be advantageous for The Ven Company to identify the customer in its pool which has acted illegally so that The Ven Company can pursue legal remedies including suing the customer for injunctive relief to prevent the unauthorised publication and/or damages, terminate its business relationship with this customer, or
30 both.

 In principle, this problem may be addressed using a robust invisible watermark scheme as defined above. However, given an image carrying a watermark, W_i , belonging to a large collection W_1, W_2, \dots, W_m , recognising which one of these marks has been used, if any, is difficult and time
35 consuming. To segregate among different customers, the number, m , of watermarks needs to be as large as the maximum number of customers that The Ven Company expects for its best selling digital images. Due to the

statistical nature of robust watermark detection, detection can be unreliable for a large value of m , and careless use of watermarks could lead to a false watermark identification, potentially leading to embarrassing problems with non-offending customers. On the other hand, the problem of a false negative is, in principle, solved by the intrinsic qualities of the watermarking scheme used, and therefore reduces the problem to selecting which mark is used out of a large collection.

A well known kind of attack on watermarks in a situation as described above is collusion of several customers. This can be described as follows: assume customers $C_{i_1}, C_{i_2}, \dots, C_{i_n}$ all buy copies of the same image I , which have been watermarked with different marks $W_{i_1}, W_{i_2}, \dots, W_{i_n}$, so that they receive image files $I_{i_1}, I_{i_2}, \dots, I_{i_n}$, respectively. These customers can together create a new image file I_c (where c stands for collusion) where each pixel (h,v) is defined by a combination of averaging and random choices among the (h,v) pixels of all images $I_{i_1}, I_{i_2}, \dots, I_{i_n}$. It has been recognised that arbitrary collusions will defeat watermarking. Prior work can be found on protection against collusion attacks, such as "Collusion secure fingerprinting for digital data" by D. Boneh, and J. Shaw, *IEEE Transactions on Information Theory*, Vol 44, No. 5, pp. 1897-1905, (1998).

It is an object of the present invention to mitigate at least some of the problems of the prior art.

Accordingly, a first aspect of the present invention provides a method for performing watermark detection, first and second collections of masks are formed, a mask from each collection is chosen to form a pair of masks, and the pair of masks are then encoded into the digital image. Next, a digital image suspected of being an unauthorised version of the encoded digital image is obtained. The suspected image is then compared with the masks in the first and second collections on a mask-by-mask basis. If the first and second masks of the pair are found in the image, the customer who distributed the unauthorised version of the encoded digital image may easily be determined, since each pair of masks has beforehand been assigned to one and only one customer. The masks in the first and second collections may be formed according to a random process. Further, to prevent a false identification of a suspected unauthorised image, the masks in the first and second collections are preferably different as determined by a predetermined criteria, eg at least one third of the values in each mask preferably should be different from the values in the other masks.

Various embodiments of the present invention provide at least one or more of the following advantages: to allow, using a method for encoding a watermark into a digital image, a person who has engaged in an unauthorised use of the digital image to be identified with greater accuracy compared with conventional methods; to support improved detection detecting of an unauthorised digital image that has been formed as a result of collusion between or among customers of a digital image vendor and to identify the colluding parties with greater accuracy as compared to conventional methods; to reduce the number of processing steps required to identify a customer or customers who engaged in an unauthorised use of a digital image; to perform collusion detection based on an analysis of several images sold to one or more customers. Various combinations of the above advantages can be realised by using a digital watermarking method which encodes two watermarks, or masks, into a digital image offered for use by a vendor. Such a method advantageously be applied for watermark detection alone, or may also be used to perform collusion detection among two or more customers. The type of watermarks and the manner in which the watermarks are encoded are determined accordingly.

To reduce the number of comparisons required to detect the watermarks in a suspected unauthorised image, the number of masks in both collections are preferably the same. As a result, embodiments can perform recognition from among N customers while performing at most $2\sqrt{N}$ tests for the case when there is no collusion, and recognises collusion of two or more customers, where \sqrt{N} stands for smallest integer larger or equal to the square root of N .

According to a second embodiment there is provided a method for collusion detection in which first and second collections of masks each include a sub-collection of masks of a first type and a sub-collection of masks of a second type. A plurality of pairs of masks are then assigned to each of a plurality of customers, with each pair of masks assigned per customer corresponds to a different one of a plurality of digital images I_i offered for use by the vendor. The first mask in each pair is derived from the first collection and the second mask is derived from the second collection, with the first mask and the second mask either both being the first type or both being of the second type. The pair of masks is then encoded into a digital image authorised for use by the associated customer.

The first and second collections of masks are formed in accordance with steps that include: a) defining at least one initial mask in a temporary collection, b) choosing a plurality of patterns for the digital

images, each of the patterns chosen for a subset of pixels in a respective one of the digital images, c) generating a first tentative mask as a possible mask to be added to the temporary collection; d) modifying the first tentative mask to carry one of said patterns $Pat(j)$ for a subset S_j of pixels (h,v) within digital image I_j , e) generating a second tentative mask which is the same as the first tentative mask except that the second tentative mask carries a pattern which is an inversion $Pat'(j)$ of the pattern $Pat(j)$ carried by the first tentative mask, f) comparing the first tentative mask and the second tentative mask to the masks already in the temporary collection, g) rejecting the first tentative mask and the second tentative mask if a predetermined correspondence exists between at least one of the first tentative mask and the second tentative mask and any of the masks in the temporary collection, and h) adding the first tentative mask and the second tentative mask to the temporary collection if the predetermined correspondence does not exist between both the first tentative mask and the second tentative mask and any of the masks in the temporary collection.

Steps c) - h) are then repeated until a desired number of masks are included in the temporary collection. The temporary collection is then divided into two groups corresponding, respectively, to the first collection and the second collection, where the sub-collection of the first type in the first collection and the second collection includes masks which carry patterns $Pat(j)$ and the sub-collection of the second type includes masks which carry patterns $Pat'(j)$.

To determine the customers who colluded to form the suspected unauthorised image, one of the following symbolic sequences $sl(j, C_i)$ for each of the vendor's customers:

$sl(j, C_i) = 0$ if C_i has not bought image I_j ,

$sl(j, C_i) = 1$ if Image I_j sold to C_i has watermarks of type A,

$sl(j, C_i) = 2$ if Image I_j sold to C_i has watermarks of type B.

An output S_j is computed for all possible pairs of customers based on the symbolic sequences. The output S_j for all possible pairs of customers are compared to the suspected image, and the parties who colluded are determined based on an outcome of the comparing step.

Another embodiment of the present method for performing collusion detection is similar to the foregoing embodiment except that instead of

using sub-collections, masks are formed for respective customers. Each mask (M) corresponds to a two-dimensional concatenation of two smaller masks. The second smaller mask (M_s') is an inversion of the first smaller mask (M_s). Further, a correlation is drawn between the first mask and a first letter of an alphabet, as well as between the second mask and a second letter of the alphabet. Mask (M) is then regarded as a code word written in the first and second letters.

According to another embodiment of the present method for performing watermark detection based on more than one image, n collections of watermarks are formed, a set of n watermarks $W(i)$ are assigned to each customer with the n watermarks being chosen from the n collections of watermarks, respectively, a different one of the watermarks from the set $W(i)$ is encoded into respective images, and the images are sold to a single customer. A group of images suspected of being unauthorised versions of the images used by the customer are then obtained, a predetermined number in the group are analysed, and the group of images is identified as being derived from the customer when watermarks are found in the predetermined number of images in the group.

According to another embodiment of the present method for performing collusion detection, n collections of watermarks are formed, watermarks are randomly chosen from the n collections, and the randomly chosen watermarks are encoded into respective images of a plurality of images. At least two customers are then authorised to use the images, with at least a portion of the images being encoded with a common watermark. A group of suspected images is then obtained, analysed to determine which images in the group have the common watermark, and parties to the collusion are determined based on the images containing the common watermark.

Further, any of the methods of the present invention may advantageously be embodied in a computer program adapted for execution on either a general purpose or special purpose computer.

Embodiments of the present invention will now be described by way of example only with reference to the accompanying drawings in which:

figure 1 shows a flow diagram illustrating a method to construct collections $Col(1)$ and $Col(2)$;

figure 2 illustrates a flow diagram showing how sub-collections A and B are constructed in accordance with an embodiment;

figure 3 depicts a flow diagram showing how symbolic sequences, $s_1(j, C_i)$, are constructed for customers in accordance with the embodiment; and

figure 4 is a flow diagram showing how checks may be performed for recognising the manner in which an image has been watermarked in accordance with an embodiment.

In accordance with a first embodiment, digital watermarks are chosen and attributed to customers based on principles of coding theory. Let $i=1, 2, \dots, N$ be an index designating customer C_i and $j=1, 2, \dots$ be an index designating picture I_j before watermarking. The present method marks an image with two watermarks and can recover both marks, when dealing with image I_j , by associating to each customer C_i a pair of masks given by:

$$(M(K(i,j,1)), M(K(i,j,2))),$$

where $M(K(i,j,1))$ belongs to some collection $Col(1)$ which may or not depend on j , and $M(K(i,j,2))$ belongs to some, preferably disjoint, collection $Col(2)$ which also, independently, may or not depend on j .

Since a square encloses the largest area of any rectangle for a given perimeter, it is more economical and thus preferable to choose $Col(1)$ and $Col(2)$ as having the same number, $Elem$, of elements. Then, there are $N' = Elem^2$ possible pairs (assuming $Col(1)$ and $Col(2)$ are disjoint collections) and the method can discriminate between N' possible customers. For simplicity, it is assumed that $N = N'$ in the sequence with no significant loss of generality. Hence, to recognise a customer, the method will only have to perform at most $2Elem$ trials to recognise the pair of masks $(M(K(i,j,1)), M(K(i,j,2)))$. By using coding theory to determine the collections $Col(1)$ and $Col(2)$, the present invention also provides protection against collusion attack, in the manner described in the discussion which follows.

Referring now to the drawings, and more particularly to figure 1, the collections $Col(1)$ and $Col(2)$ are formed by constructing masks using a random process and rejecting any mask too close to the masks already selected. More precisely, a random number generator 101 is used to generate an array of 0's and 1's at 102. Using some preferred method to generate masks at 103, such as the one described in US Patent No. 5,825,892 issued to Mintzer and Braudaway, the content of which is hereby incorporated by reference, a tentative mask is generated at 104. Using a comparator 105, the tentative mask 104 is compared to all masks already selected in the union of $Col(1)$ and $Col(2)$ constructed so far in a temporary collection

106. If the tentative mask 104 is too close to at least one mask of temporary collection 106, the mask is rejected. Otherwise, the tentative mask is added to collection 106, and the overall process stops when N masks have been selected. Thus, coding theory will first be used to choose all elements of the union of Col(1) and Col(2) to be sufficiently different. For instance, it is easy to ensure that any two elements chosen from these two collections will differ on at least one third of their entries. Such consideration will be preferable to avoid false identification of images as being unauthorised. It will be appreciated that the degree of acceptable correspondence between an existing mask and the tentative mask 104 will vary according to application as will the type of measure of correspondence between an existing mask and a tentative mask 104.

When composing collections Col(1) and Col(2), which can be done by random trials and rejection of bad cases, a further restriction can be imposed which will serve to detect images produced as a result of a collusion between two customers. To that effect, for each j , a subset S_j of the values of (h,v) is chosen. A pattern $Pat(j)$ of "0's" and "1's" is chosen. A pattern $Pat'(j)$ is obtained by changing each entry of $Pat(j)$ to 1 minus this entry (so that 1 turns to 0 and 0 turns to 1). Half of the elements of Col(1) and Col(2) will carry $Pat(j)$ on S_j , the other halves carrying $Pat'(j)$.

Thus, both Col(1) and Col(2) are cut in two halves, say Col(1) decomposes as the disjoint union of Col(1)A and Col(1)B, while Col(2) decomposes as the disjoint union of Col(2)A and Col(2)B. It will be appreciated that A corresponds to $Pat(j)$ being carried on S_j , while B corresponds to $Pat'(j)$ being carried on S_j . A customer, C_i , is assigned a pair of watermarks both of type A, or both of type B.

With reference to figure 2, there is illustrated an embodiment of a process for constructing sub-collections A and B. All steps are identical to corresponding steps illustrated in figure 1, except that after one tentative mask has been constructed at 104 for A, a second mask is constructed at 204 for B, by only inverting the elements of S_j as described above. Both masks (for A and for B) are simultaneously accepted or rejected in the comparator 105. When two customers collude, the S_j of the image they produce by collusion carries:

$Pat(j)$ if they are both in A

$Pat'(j)$ if they are both in B

Neither of the above otherwise.

Figure 3 shows a flowchart for the construction of a symbolic sequence $s_1(j, C_i)$ for each customer, by setting:

$s_1(j, C_i) = 0$ if C_i has not bought image I_j ,

$s_1(j, C_i) = 1$ if Image I_j sold to C_i has watermarks of type A,

$s_1(j, C_i) = 2$ if Image I_j sold to C_i has watermarks of type B.

Once such sequences have been computed, the output on S_j of each pair of customers can be computed, and even of more complex combinations of them. In case only one pair of customers collude for some long series of images (which need not be consecutive), the pair will be easily discovered, assuming the A-B partitions are made so as to generate $s_1(j, C_i)$ which diverges promptly from one customer to the next one as j increases. For instance, a way to organise this partition as j evolves is depicted in Figure 1.

If more than one pair of customers colludes, the recognition is more complex. However, the list of suspects can be restricted and the partitions can be re-organised so as to catch the wrongdoers. From classical coding theory, it is evident that the restriction on marks imposed by limiting the freedom on S_j leaves a large number of codes apart from each other. Let P denote the number of entries of the matrices M (where M stands for a mask such as $M(K(i, j, 1))$ or $M(K(i, j, 2))$). Recall that S_j contains some number of entries of the matrices M , say a third, i.e., the size of S_j is $P/3$. As we deal now with rough estimates, assume that P is a multiple of 6 and set $P = 3P' = 6P''$. Using these notations, it was previously suggested that different codes in the union of $\text{Col}(1)$ and $\text{Col}(2)$ differs in at least P' bits. The number of binary codes of length L which are separated by a distance d is usually denoted by $A(L, d)$. The number of binary codes of length L with Q "1's" which are separated by a distance d is usually denoted by $A(L, d, Q)$: both concepts are discussed at length in the coding literature (see, e.g., "The Theory of Error Correcting Codes", North-Holland Mathematical Library, V.16, by Florence J. Macwilliams and Neil J. A. Sloane, North-Holland, Amsterdam (1977)). Under the constraints imposed, the freedom of choice for matrices M is measured as $A(3P', P')$. $A(3P', P')$ is quite large when P' is of the order of hundreds as in any high resolution image application (more generally, $A(\alpha P', P')$

grows exponentially with respect to P' when $\alpha > 2$). Even if about half of the entries of the matrix out of S_j are "1's", the number of possible choices is quite large as it is given by $A((\alpha P', P', P'))$ for α close to 2, see, for example, Chapter 17 and Appendix A of the previously mentioned book by MacWilliams and Sloane for lower bound estimates of these numbers which confirm these statements.

As disclosed previously, using two watermarks from a pair of collections of watermarks reduces considerably the number of checks which have to be made on average to recognise how an image has been watermarked. More precisely, as shown in figure 4, at most only $2 \times \text{Elem}$ tests need to be made: the masks are selected one by one from the two collections $\text{Col}(1)$ and $\text{Col}(2)$ at 106, each containing Elem masks. The selected mask is used in step 401 to construct a watermark 402 according to the watermarking method. The presence of this watermark in the image 403 is then tested in decision block 404. Clearly, more than two marks put together would shorten even more the verification process.

An alternate embodiment of the method of the present invention uses some of the concepts discussed above to watermark differently several copies of a given image. The basic principle that different marks should correspond to masks, M , having at least a predetermined distance or difference should, preferably, be observed, and the use of the triple $(S_j, \text{Pat}(j), \text{Pat}'(j))$ can still be employed. In fact, these principles together allow for reasonable results. The embodiment will use a specific way to generate different codes so that these codes are still quite different if restricted to a reasonably large portion of the image. The method can also be used for composing the collections $\text{Col}(1)$ and $\text{Col}(2)$ discussed previously.

According to this method, each mask, M , will be conceived as a two-dimensional concatenation of smaller matrices. A single small matrix, M_s , together with the matrix $M_s' = 1 - M_s$ (so that "0's" become "1's", and "1's" become "0's") can be generated. M_s and M_s' can be considered as two letters of an alphabet, and M as a code word written with these letters. The same coding theory considerations now apply to generate a proper collection of different masks from the constituents M_s and M_s' . In the case of this method being used to generate $\text{Col}(1)$ and $\text{Col}(2)$, the same pair may or may not be used for both $\text{Col}(1)$ and $\text{Col}(2)$. The analysis is mostly performed on a single image I .

Next, consider the case where the detection occurs based on several images. In this case, the determination of customer identity is performed by analysing several images sold to the same customer and collusion detection is performed by analysing several corrupted images resulting from the collusion of the same subset of customers. For instance, to identify customers, each customer C_i will be assigned a set of n watermarks $W(i)$, one from each of n collections $Col(1), \dots, Col(n)$. The images sold to C_i will each carry one watermark from the set $W(i)$. By analysis of the watermarks on images sold to C_i , the identity of C_i can be determined. One way to get the set $W(i)$ is to choose randomly n watermarks, one from each of the collections $Col(1), \dots, Col(n)$ and to keep track of the set $W(i)$ for each customer C_i .

In accordance with another embodiment of the present invention, a similar technique is used to detect collusion. For each image I_j sold to a customer, C_i , a watermark is chosen randomly from n collections of watermarks $Col(1), \dots, Col(n)$ and applied to the image. Preferably, n is chosen to be smaller than the number of customers. Further, the owner needs to keep track of which watermark is applied to an image, I_j , sold to a customer, C_i , in a list L . When two customers C_1 and C_2 collude to attempt to remove the watermark and generate a set of modified images, some images will have the same watermark applied to images of both C_1 and C_2 which thus cannot be removed by the collusion attack, as described earlier. By finding the images that have a common watermark, extracting these common watermarks and comparing the results with the list L , it can be determined which two customers have been colluding. This idea can be used to detect collusion by more than two customers.

The method described here for images can apply as well to videos or other forms of animated images, or to sounds, by anyone versed in the art. The above described methods may be implemented on a suitably programmed general purpose or specific computer system.

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognise that the invention can be practiced with modification within the spirit and scope of the appended claims.

CLAIMS

1. A digital watermarking method, comprising:

forming a first collection of masks;

forming a second collection of masks, each mask in the second collection being different from the masks in the first collection;

assigning to each of a plurality of customers C_i a plurality of pairs of masks $(M(K(i,j,1)), M(K(i,j,2)))$, each pair of masks assigned corresponding to a different image of a plurality of digital images I_j , wherein a first mask $M(K(i,j,1))$ in each pair is derived from the first collection and a second mask $M(K(i,j,2))$ in each pair is derived from the second collection; and

encoding into a specific digital image of the plurality of digital images I_j a pair of masks assigned to one of the customers which corresponds to the specific digital image.

2. A method as claimed in claim 1, further comprising:

obtaining a digital image suspected of being an unauthorised version of the specific digital image;

comparing the masks in the first collection to the suspected digital image on a mask-by-mask basis;

identifying a mask in the first collection which is found in the suspected mask;

comparing the masks in the second collection to the suspected digital image on a mask-by-mask basis;

identifying a mask in the second collection which is found in the suspected mask;

forming the mask identified in the first collection and the mask identified in the second collection into a pair of identified masks;

determining which of the plurality of customers was assigned to the pair of identified masks for the specific digital image; and

identifying the customer assigned to the pair of identified masks for the specific digital image as being responsible for the unauthorised version of the specific digital image.

3. A method as claimed in claim 2 in which the comparison comprises the step of

detecting whether a watermark generated from a mask is embedded in the specific digital image.

4. A method as claimed in any preceding claim, in which the steps of forming the first collection and the second collection include:

- a) defining at least one initial mask in a temporary collection;
- b) generating a tentative mask as a possible mask to be added to the temporary collection;
- c) comparing the tentative mask to the masks already in the temporary collection;
- d) rejecting the tentative mask if a predetermined correspondence exists between the tentative mask and any of the masks in the temporary collection;
- e) adding the tentative mask to the temporary collection if the predetermined correspondence does not exist;
- f) repeating steps b) - e) until a desired number of masks are included in the temporary collection; and
- g) dividing the desired number of masks in the temporary collection into two groups corresponding, respectively, to the first collection and the second collection.

5. A method as claimed in claim 4, in which the predetermined correspondence is such that less than one third of values in the tentative mask are different to values in any of the masks in the temporary collection.

6. A method as claimed in either of claims 4 and 5, in which step b) includes the steps of randomly generating an array of zeros and ones; and generating the tentative mask from the array.

7. A method as claimed in any preceding claim in which the first collection and the second collection contain a same number of masks.

8. A method as claimed in any preceding claim in which the first collection of masks and the second collection of masks are disjoint collections having a same number (Elem) of masks, and wherein at most $2 \times$ (Elem) comparisons are performed to determine which of the plurality of customers was assigned to the pair of identified masks for the specific digital image.

9. A method as claimed in any preceding claim in which the first collection of masks includes a sub-collection of a first type of mask and a sub-collection of a second type of mask, the second collection of masks includes a sub-collection of a first type of mask and a sub-collection of a

second type of mask wherein the first mask and the second mask in each pair either both being of the first type of mask or both being of the second type of mask.

10. A method as claimed in claim 9 in which the steps of forming the first collection and the second collection include the steps of:

(a) defining at least one initial mask in a temporary collection;

(b) choosing a plurality of patterns for the digital images, each of the patterns being chosen for a subset of pixels in a respective one of the digital images;

(c) generating a first tentative mask as a possible mask to be added to the temporary collection;

(d) modifying the first tentative mask to carry one of the patterns $Pat(j)$ for a subset S_j of pixels (h,v) within digital image I_j ;

(e) generating a second tentative mask, the second tentative mask being same as the first tentative mask except that the second tentative mask carries a pattern which is an inversion $Pat'(j)$ of the pattern $Pat(j)$ carried by the first tentative mask;

(f) comparing the first tentative mask and the second tentative mask to the masks already in the temporary collection;

(g) rejecting the first tentative mask and the second tentative mask if a predetermined correspondence exists between at least one of the first tentative mask and the second tentative mask and any of the masks in the temporary collection;

(h) adding the first tentative mask and the second tentative mask to the temporary collection if the predetermined correspondence does not exist between both the first tentative mask and the second tentative mask and any of the masks in the temporary collection;

(i) repeating steps c) - h) until a desired number of masks are included in the temporary collection; and

(j) dividing the desired number of masks in the temporary collection into two groups corresponding, respectively, to the first collection and the second collection, wherein the sub-collection of the first type in the first collection and the second collection includes masks which carry patterns $Pat(j)$ and the sub-collection of the second type includes masks which carry patterns $Pat'(j)$.

11. A method as claimed in any preceding claim, further comprising the step of constructing one of the following symbolic sequences $s_1(j, C_i)$ for each of the plurality of customers:

$s_1(j, C_i) = 0$ if C_i has not bought image I_j ,

$s1(j, C_i) = 1$ if Image I_j sold to C_i has watermarks of type A, $s1(j, C_i) = 2$ if Image I_j sold to C_i has watermarks of type B,

computing an output S_j for all possible pairs of the plurality of customers based on the symbolic sequences assigned to the customers in the constructing step;

obtaining a digital image suspected of being an unauthorised version of the specific digital image;

comparing the output S_j for all of the possible pairs of customers as determined in the computing step to the suspected digital image; and

identifying which two customers colluded to produce the unauthorised version of the specific digital image based on an outcome of the comparing step.

12. A digital watermarking system, comprising:

means for forming a first collection of masks;

means for forming a second collection of masks, each mask in the second collection being different from the masks in the first collection;

means for assigning to each of a plurality of customers C_i a plurality of pairs of masks $(M(K(i,j,1)), M(K(i,j,2)))$, each pair of masks assigned corresponding to a different image of a plurality of digital images I_j , wherein a first mask $M(K(i,j,1))$ in each pair is derived from the first collection and a second mask $M(K(i,j,2))$ in each pair is derived from the second collection; and

means for encoding into a specific digital image of the plurality of digital images I_j a pair of masks assigned to one of the customers which corresponds to the specific digital image.

13. A system as claimed in claim 12, further comprising:

means for obtaining a digital image suspected of being an unauthorised version of the specific digital image;

means for comparing the masks in the first collection to the suspected digital image on a mask-by-mask basis;

means for identifying a mask in the first collection which is found in the suspected mask;

means for comparing the masks in the second collection to the suspected digital image on a mask-by-mask basis;

means for identifying a mask in the second collection which is found in the suspected mask;

means for forming the mask identified in the first collection and the mask identified in the second collection into a pair of identified masks;

means for determining which of the plurality of customers was assigned to the pair of identified masks for the specific digital image; and

means for identifying the customer assigned to the pair of identified masks for the specific digital image as being responsible for the unauthorised version of the specific digital image.

14. A system as claimed in claim 13 in which the means for comparing comprises:

means for detecting whether a watermark generated from a mask is embedded in the specific digital image.

15. A system as claimed in any of claims 12 to 14, in which the means for forming the first collection and the second collection includes:

means for defining at least one initial mask in a temporary collection;

means for generating a tentative mask as a possible mask to be added to the temporary collection;

means for comparing the tentative mask to the masks already in the temporary collection;

means for rejecting the tentative mask if a predetermined correspondence exists between the tentative mask and any of the masks in the temporary collection;

means for adding the tentative mask to the temporary collection if the predetermined correspondence does not exist;

means for repeating the above means until a desired number of masks are included in the temporary collection; and

means for dividing the desired number of masks in the temporary collection into two groups corresponding, respectively, to the first collection and the second collection.

16. A system as claimed in claim 15, in which the predetermined correspondence is such that less than one third of values in the tentative mask are different to values in any of the masks in the temporary collection.

17. A system as claimed in either of claims 15 and 16, in which the means for generating a tentative mask includes means for randomly generating an array of zeros and ones; and generating the tentative mask from the array.

18. A system as claimed in any of claims 12 to 17 in which the first collection and the second collection contain a same number of masks.

19. A system as claimed in any of claims 12 to 18 in which the first collection of masks and the second collection of masks are disjoint collections having a same number (Elem) of masks, and wherein at most $2 \times$ (Elem) comparisons are performed to determine which of the plurality of customers was assigned to the pair of identified masks for the specific digital image.

20. A system as claimed in any of claims 12 to 19 in which the first collection of masks includes a sub-collection of a first type of mask and a sub-collection of a second type of mask, the second collection of masks includes a sub-collection of a first type of mask and a sub-collection of a second type of mask wherein the first mask and the second mask in each pair either both being of the first type of mask or both being of the second type of mask.

21. A system as claimed in claim 20 in which the means for forming the first collection and the second collection includes

means for defining at least one initial mask in a temporary collection;

means for choosing a plurality of patterns for the digital images, each of the patterns being chosen for a subset of pixels in a respective one of the digital images;

means for generating a first tentative mask as a possible mask to be added to the temporary collection;

means for modifying the first tentative mask to carry one of the patterns Pat(j) for a subset S_j of pixels (h,v) within digital image I_j ;

means for generating a second tentative mask, the second tentative mask being same as the first tentative mask except that the second tentative mask carries a pattern which is an inversion Pat'(j) of the pattern Pat(j) carried by the first tentative mask;

means for comparing the first tentative mask and the second tentative mask to the masks already in the temporary collection;

means for rejecting the first tentative mask and the second tentative mask if a predetermined correspondence exists between at least one of the first tentative mask and the second tentative mask and any of the masks in the temporary collection;

means for adding the first tentative mask and the second tentative mask to the temporary collection if the predetermined correspondence does

not exist between both the first tentative mask and the second tentative mask and any of the masks in the temporary collection;

means for repeating the above means until a desired number of masks are included in the temporary collection; and

5 means for dividing the desired number of masks in the temporary collection into two groups corresponding, respectively, to the first collection and the second collection, wherein the sub-collection of the first type in the first collection and the second collection includes masks which carry patterns $Pat(j)$ and the sub-collection of the second type
10 includes masks which carry patterns $Pat'(j)$.

22. A system as claimed in any of claims 12 to 21, further comprising means for constructing one of the following symbolic sequences $s1(j, C_i)$ for each of the plurality of customers:

15 $s1(j, C_i) = 0$ if C_i has not bought image I_j , $s1(j, C_i) = 1$ if Image I_j sold to C_i has watermarks of type A,

$s1(j, C_i) = 2$ if Image I_j sold to C_i has watermarks of type B,

means for computing an output S_j for all possible pairs of the plurality of customers based on the symbolic sequences assigned to the
20 customers by the means for constructing;

means for obtaining a digital image suspected of being an unauthorised version of the specific digital image;

means for comparing the output S_j for all of the possible pairs of customers as determined in the computing step to the suspected digital
25 image; and

means for identifying which two customers colluded to produce the unauthorised version of the specific digital image based on an outcome of the means for comparing.

30 23. A digital watermarking method substantially as described herein with reference to and/or as illustrated in the accompanying drawings.

24. A digital watermarking system substantially as described herein with reference to and/or as illustrated in the accompanying drawings.



Application No: GB 0029063.5
Claims searched: 1-24

Examiner: Frank D. Moeschler
Date of search: 3 May 2001

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK Cl (Ed.S): H4F (FBB)
Int Cl (Ed.7): G06T-1/00; H04N-1/00; 1/32, 5/913
Other: Online: WPI; JAPIO ;EPODOC; INSPEC; IEEEExplore

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	WO 99/10858 A2 (LEIGHTON) See whole document - especially pages 10-15	1-3, 12-14
X	WO 97/34391 A1 (LEIGHTON) See whole document - especially pages 12-14	1-3, 12-14
X	EP 0933919 A2 (CANON) See whole document - especially Fig 2	1-3, 12-14
A	IEEE trans. Multimedia, Vol 2 No 4, Dec 2000, C-S Lu, S-K Huang, C-J Sze, H-Y M Liao, "Cocktail Watermarking...", pp 209-224, especially pp212.	

X Document indicating lack of novelty or inventive step
Y Document indicating lack of inventive step if combined with one or more other documents of same category.
& Member of the same patent family

A Document indicating technological background and/or state of the art.
P Document published on or after the declared priority date but before the filing date of this invention.
E Patent document published on or after, but with priority date earlier than, the filing date of this application.